

**PCT**WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau

## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<b>(51) International Patent Classification <sup>6</sup> :</b> <b>G06F 17/30</b>	<b>A1</b>	<b>(11) International Publication Number:</b> <b>WO 99/22315</b> <b>(43) International Publication Date:</b> 6 May 1999 (06.05.99)
<b>(21) International Application Number:</b> PCT/US98/20719 <b>(22) International Filing Date:</b> 2 October 1998 (02.10.98)  <b>(30) Priority Data:</b> 08/959,058                      28 October 1997 (28.10.97)                      US  <b>(71) Applicant:</b> CACHEFLOW, INC. [US/US]; 650 Almanor Avenue, Sunnyvale, CA 94086 (US).  <b>(72) Inventors:</b> MALCOLM, Michael; 250 Family Farm Road, Woodside, CA 94062 (US). TELFORD, Ian; 72 John Street East, Waterloo, Ontario N2J 1G1 (US).  <b>(74) Agent:</b> SWERNOFSKY, Steven, A.; The Law Offices of Steven A. Swernofsky, P.O. Box 390013, Mountain View, CA 94039-0013 (US).		<b>(81) Designated States:</b> CA, CN, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).  <b>Published</b> <i>With international search report.</i>
<b>(54) Title:</b> ADAPTIVE ACTIVE CACHE REFRESH  <b>(57) Abstract</b>  The invention provides a system and system for automatically refreshing documents in a cache, so that each particular document is refreshed no more often and no less often than needed. For each document, the cache estimates a probability distribution of times for client requests for that document and a probability distribution of times for server changes to that document. Times for refresh are selected for each particular document in response to both the estimated probability distribution of times for client requests and the estimated probability distribution of times for server changes. The invention also provides a system and system for objectively estimating the value the cache is providing for the system including the cache. The cache estimates for each document a probability distribution of times for client requests for that document, and determines a cumulative probability distribution which reflects the estimated marginal hit rate at the storage limit of the cache and the marginal advantage of adding storage to the cache.		

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Larvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav	TM	Turkmenistan
BF	Burkina Faso	GR	Greece		Republic of Macedonia	TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

1 Title of the Invention

2

3 Adaptive Active Cache Refresh

4

5 Background of the Invention

6

7 1. *Field of the Invention*

8

9 This invention relates to caches.

10

11 2. *Related Art*

12

13 When a client device seeks to obtain information from server devices on a net-  
14 work, it is sometimes desirable to provide a cache, that is, a device which maintains copies of  
15 that information so that multiple requests for the same information can be satisfied at the cache,  
16 and do not require that information to be transmitted repeatedly across the network. Known  
17 caches do this to reduce the amount of communication bandwidth used between the clients and  
18 the servers, and when shared by more than one client, act to reduce the total amount t of com-  
19 munication bandwidth used between all of the clients and the servers.

20

21 One problem in the art is that information requested a second time (possibly re-  
22 quested a second time by the same client, or requested by a second client after a first client has  
23 already requested that information once) can change at the server between the time it is first re-  
24 quested and the time it is requested again. In such cases, transmitting the stored information  
25 from the cache gives inaccurate information to the second requester. This can reduce the confi-  
26 dence users at the client devices have for the information provided by the cache.

27

28 One known method is to transmit each request from the client device to the  
29 server. so as to obtain an answer as to whether the document must be refreshed before it is  
30 served (transmitted) to the client device. While this method achieves the purpose of serving  
31 only refreshed information to the client device, it has the drawback that the client device must  
32 wait for contact with the server device and the reply from the server device, even when the in-  
33 formation is already present in the cache. Moreover, this method uses communication band-  
34 width by sending requests to the server device and receiving confirmations from the server de-  
35 vice which can be unnecessary.

1  
2 It would be advantageous to provide a cache which reduces the average amount  
3 of time users at the client device wait for information, rather than attempting to reduce the  
4 amount of communication bandwidth between the client devices and the server devices. One  
5 aspect of the invention is to automatically refresh the information maintained in the cache, not-  
6 withstanding that this uses additional communication bandwidth. The cache occasionally que-  
7 ries the server device to determine if the document has been changed, so the cache can maintain  
8 an up-to-date version of the document. When the document is requested by the client device,  
9 the cache serves that document immediately without checking with the server device.

10  
11 Refreshing information in the cache is useful, but some documents require refresh  
12 more often than others. If a particular document is selected for refresh less often than required,  
13 it will sometimes be served to the client device even though it is "stale" (that is, modified at the  
14 server since the cache last obtained a copy). In contrast, if a particular document is selected for  
15 refresh more often than required, the document will sometimes be refreshed unnecessarily, thus  
16 wasting communication bandwidth.

17  
18 Accordingly, it would be advantageous to provide a method and system for re-  
19 freshing documents so that each particular document is refreshed no more often and no less of-  
20 ten than needed. This advantage is achieved in a system in which the times for refresh are tai-  
21 lored to each particular document, in response to both an estimated probability distribution of  
22 times for client requests for that document and an estimated probability distribution of times for  
23 server changes to that document.

24  
25 Another problem in the art is that it is difficult to objectively determine the value  
26 the cache is providing for the system including the cache, or whether the cache is too small or  
27 too large. In contrast with persistent storage devices, for which it is easy to determine how full  
28 they are and whether the size of the storage device is too small or too large, the cache is nearly  
29 always nearly full of data being stored for later request by the client device.

30  
31 One known method to determine the value of the cache is to measure the cache  
32 "hit rate," that is, the fraction of information requests which are for documents already main-  
33 tained in the cache. However, this measure is extremely dependent on the degree of locality of  
34 reference to information requested by the client device, which in turn is extremely dependent on

1 the number of client devices, the nature of information they are requesting, and the rate at  
2 which they request that information.

3  
4 Accordingly, it would be advantageous to provide a method and system for ob-  
5 jectively estimating the value the cache is providing for the system including the cache, such as  
6 whether the cache is too small or too large for selected objectives. This advantage is achieved  
7 in a system in which the cache estimates for each document a probability distribution of times  
8 for client requests for that document, and determines a cumulative probability distribution  
9 which reflects the estimated marginal hit rate at the storage limit of the cache and the marginal  
10 advantage of adding storage to the cache.

### 11 12 Summary of the Invention

13  
14 The invention provides a method and system for automatically refreshing docu-  
15 ments in a cache, so that each particular document is refreshed no more often and no less often  
16 than needed. For each document, the cache estimates a probability distribution of times for cli-  
17 ent requests for that document and a probability distribution of times for server changes to that  
18 document. Times for refresh are selected for each particular document in response to both the  
19 estimated probability distribution of times for client requests and the estimated probability dis-  
20 tribution of times for server changes.

21  
22 The invention also provides a method and system for objectively estimating the  
23 value the cache is providing for the system including the cache. The cache estimates for each  
24 document a probability distribution of times for client requests for that document, and deter-  
25 mines a cumulative probability distribution which reflects the estimated marginal hit rate at the  
26 storage limit of the cache and the marginal advantage of adding storage to the cache.

### 27 28 Brief Description of the Drawings

29  
30 Figure 1 shows a block diagram of a system for periodically refreshing docu-  
31 ments in a cache.

32  
33 Figure 2 shows a process flow diagram of a method for periodically refreshing  
34 documents in a cache.

Detailed Description of the Preferred Embodiment

In the following description, a preferred embodiment of the invention is described with regard to preferred process steps and data structures. Those skilled in the art would recognize after perusal of this application that embodiments of the invention can be implemented using general purpose processors or special purpose processors or other circuits adapted to particular process steps and data structures described herein, and that implementation of the process steps and data structures described herein would not require undue experimentation or further invention.

Inventions disclosed herein can be used in conjunction with inventions disclosed in one or more of the following patent applications:

Provisional U.S. Application 60/048,986, filed June 9, 1997, in the name of inventors Michael Malcolm and Robert Zarnke, titled "Network Object Cache Engine." assigned to CacheFlow, Inc., attorney docket number CASH-001.

U.S. Application Serial No. 08/\_\_\_\_\_, filed this same day, in the name of inventors Doug Crow, Bert Bonkowski, Harold Czegledi, and Tim Jenks. titled "Shared Cache Parsing and Pre-fetch," assigned to CacheFlow, Inc., attorney docket number CASH-004.

These applications are referred to herein as the "Cache Disclosures." and are hereby incorporated by reference as if fully set forth herein.

*System Elements*

Figure 1 shows a block diagram of a system for periodically refreshing documents in a cache.

A system 100 includes a cache 110, at least one client device 120, and at least one server device 130. Each client device 120 is coupled to the cache 110 using a client communication path 121, such as a dial-up connection, a LAN (local area network), a WAN (wide area network), or some combination thereof. Similarly, each server device 130 is also coupled to the cache 110 using a server communication path 131, such as a dial-up connection, a LAN (local

1 area network), a WAN (wide area network), or some combination thereof. In a preferred em-  
2 bodiment, the client communication path 121 includes a LAN, while the server communication  
3 path 131 includes a network of networks such as an internet or intranet.

4  
5 As used herein, the terms "client" and "server" refer to a relationship between the  
6 client or server and the cache 110, not necessarily to particular physical devices. As used herein,  
7 one "client device" 120 or one "server device" 130 can comprise any of the following: (a) a sin-  
8 gle physical device capable of executing software which bears a client or server relationship to  
9 the cache 110; (b) a portion of a physical device, such as a software process or set of software  
10 processes capable of executing on one hardware device, which portion of the physical device  
11 bears a client or server relationship to the cache 110; or (c) a plurality of physical device, or por-  
12 tions thereof, capable of cooperating to form a logical entity which bears a client or server rela-  
13 tionship to the cache 110. The phrases "client device" 120 and "server device" 130 refer to such  
14 logical entities and not necessarily to particular individual physical devices.

15  
16 The server device 130 includes memory or storage 132 having a web document  
17 133. In a preferred embodiment, the web document 133 can include text and directions for dis-  
18 play, pictures, such as data in GIF or JPEG format, other multimedia data, such as animation,  
19 audio (such as streaming audio), movies, video (such as streaming video), program fragments,  
20 such as Java, Javascript, or ActiveX, or other web documents, such as when using frames.

21  
22 The cache 110 includes a processor 111, program and data memory 112, and  
23 mass storage 113. The cache 110 maintains a first set of web objects 114 in the memory 112 and  
24 a second set of web objects 114 in the storage 113.

25  
26 In a preferred embodiment, the cache 110 includes a cache device such as de-  
27 scribed in the Cache Disclosures defined herein, hereby incorporated by reference as if fully set  
28 forth therein.

29  
30 The cache 110 receives requests from the client device 120 for a web object 114  
31 and determines if that web object 114 is present at the cache 110, either in the memory 112 or in  
32 the storage 113. If the web object 114 is present in the memory 112, the cache 110 transmits the  
33 web object 114 to the client device 120 using the client communication path 121. If the web  
34 object 114 is present in the storage 113 but not in the memory 112, the cache 110 loads the web  
35 object 114 into the memory 112 from the storage 113, and proceeds as in the case when the web

object 114 was originally present in the memory 112. If the web object 114 is not present in either the memory 112 or the storage 113, the cache 110 retrieves the web object 114 from the appropriate server device 130, places the web object 114 in the memory 112 and the storage 113, and proceeds as in the case when the web object 114 was originally present in the memory 112.

Due to the principle of locality of reference, it is expected that the cache 110 will achieve a substantial "hit rate," in which many requests from the client device 120 for web objects 114 will be for those web objects 114 already maintained by the cache 110, reducing the need for requests to the server device 130 using the server communication path 131.

#### *Determining Expected Frequencies*

For each web object 114, the cache 110 determines a probability that the web object 114 is stale at time  $t$ , and a probability that the web object 114 is requested at request  $h$ :

$$\text{Psi}(t) = \text{Probability} \{ \text{object } i \text{ is stale at time } t \} \quad (141)$$

$$\text{Pri}(h) = \text{Probability} \{ \text{object } i \text{ is requested at request } h \} \quad (142)$$

The probability of request  $\text{Pri}(h)$  is measured for "request  $h$ " rather than for "time  $t$ ," so that the probability  $\text{Pri}(h)$  is not distorted by changes in the frequency of requests. The frequency of requests can change, for example, while the cache 110 is being moved from one location to another, such as its original shipment, or while there are relatively few client devices 120 requesting web objects 114 from the cache 110, such as early in the morning or late at night. The probability  $\text{Pri}(h)$  is thus responsive to "request time"  $h$  rather than "wall-clock time"  $t$ .

In a preferred embodiment, the cache 110 maintains a value for "request time"  $h$ , since a request-time of zero when the cache 110 was first manufactured, and stores that value in a non-volatile memory.

Having defined the probabilities  $\text{Psi}(t)$  and  $\text{Pri}(h)$ , the probability that the web object 114 will be served stale by the cache 110 on the next request is the product of the probabilities  $\text{Psi}(t) \cdot \text{Pri}(h)$ .

$$\text{Pi}(\text{current time, current request}) = \text{Psi}(\text{current time}) \cdot \text{Pri}(\text{current request})$$



1                   = Probability { object i is stale at this time and  
 2                   object i is requested at this request-time }  
 3           (143)

4  
 5                   Thus, each web object 114 i has a corresponding product  $P_i$  (current time, current  
 6 request-time), which indicates the probability that the web object 114 will be served stale by the  
 7 cache at the next request. The sum of such products  $P_i$  (current time, current request-time) for  
 8 all web objects 114 i in the cache 110 is the cumulative probability that the next web object 114  
 9 requested by one of the client devices 120 will be served stale by the cache 110.

10

11                   The cache 110 chooses to attempt to refresh the web object 114 with the highest  
 12 such product  $P_i$  (current time, current request-time). The cache 110 automatically attempts to  
 13 refresh web objects 114 until the cumulative probability of all products  $P_i$  (current time, current  
 14 request-time) is less than a selected threshold value. In a preferred embodiment, the selected  
 15 threshold value is between about 1% and about 5%.

16

17                   The probabilities  $P_{si}(t)$  and  $P_{ri}(h)$  each follow a Poisson distribution, that is, that  
 18 the probability of a particular web object 114 becoming stale in any particular time interval is  
 19 responsive to a random process having a Poisson distribution, the probability of a particular web  
 20 object 114 being requested in any particular request-time interval is also responsive to a random  
 21 process having a Poisson distribution.

22

23                   Accordingly, the cache 110 estimates  $P_{si}(t)$  as follows:

24

$$25 \quad P_{si}(t) = 1 - \exp(-a t) \quad (144)$$

26

27                   where the function **exp** is exponentiation using the natural base e, and  
 28 the value a is a parameter of the Poisson process.

29

30                   It follows that

31

$$32 \quad A = \ln(2) / EUI \quad (145)$$

33

34                   where  $\ln(2)$  is the natural logarithm of 2 (approximately 0.69315), and  
 35 EUI is the estimated mean interval between updates to the web object 114 i at the

1 server device 130.

2

3 EUI and similar values described herein are specific to each web object 114 i, and  
4 are determined separately for each web object 114 i. However, for convenience in notation, the  
5 term "EUI" and similar terms are not subscripted to so indicate.

6

7 Accordingly, the cache 110 estimates EUI in response to the times of actual up-  
8 dates of the web object 114 at the server device 130, and is able to determine  $\Psi(t)$  in response  
9 to EUI.

10

11 Similarly, the cache 110 estimates  $P_{ri}(h)$  as follows:

12

13 
$$P_{ri}(h) = b \exp(-b h) \quad (146)$$

14

15 where

16 the value  $b$  is a parameter of the Poisson process.

17

18 It follows that

19

20 
$$P_{ri}(\text{current request-time}) = \ln(2) / EAI \quad (147)$$

21 where EAI is the estimated mean interval between requests for the web object  
22 114 i at the cache 110 (that is, from any client device 120).

23

24 Accordingly, the cache 110 estimates EAI in response to the request-times of ac-  
25 tual requests for the web object 114 from any client device 120, and is able to determine  $P_{ri}(h)$   
26 in response to EAI.

27

28 *Method of Operation*

29

30 Figure 2 shows a process flow diagram of a method for periodically refreshing  
31 documents in a cache.

32

33 A method 200 includes a set of flow points to be noted, and steps to be executed,  
34 cooperatively by the system 100, including the cache 110, the client device 120, and the server  
35 device 130.

1  
2 At a flow point 210, a particular web object 114 is loaded into the cache 110.  
3 The particular web object 114 can be loaded into the cache 110 in one of two ways:

4  
5 Initial load: The web object 114 was first requested from the cache 110 by one of the  
6 client devices 120, found not to be maintained in the cache 110, requested by the cache  
7 110 from one of the server devices 130, transmitted by the server device 130 to the cache  
8 110, stored by the cache 110, and transmitted by the cache 110 to the requesting client  
9 device 120; or

10  
11 Reload: The web object 114 was maintained in the cache 110 after a previous initial load  
12 or reload, found by the cache 110 to be stale, and reloaded by the cache 110 from one of  
13 the server devices 130.

14  
15 At a step 211, the cache 110 makes an initial estimate of the values of EUI and  
16 EAI for the particular web object 114.

17  
18 In a preferred embodiment, the cache 110 performs the step 211 by determining a  
19 type for the web object 114, and making its initial estimate of the values of EUI and EAI in re-  
20 sponse to that type. The cache 110 determines the type for the web object 114 in response to  
21 information in the HTTP (hypertext transfer protocol) response made by the server device 130.  
22 The web object 114 can be one of the following types:

23  
24 type T: a web object 114 for which an expiration time (XT), sometimes called a "time to  
25 live" or "TTL," is specified by the server device 130.

26  
27 type M: a web object 114 for which a last modified time (MT) is specified, but no XT is  
28 specified, by the server device 130;

29  
30 type N: a web object 114 for which no MT or XT is specified by the server device 130.

31  
32 In a preferred embodiment, the cache 110 estimates a value for EUI in response  
33 to the type determined for the web object 114.

For the initial load of the web object 114 into the cache 110, the following values are used:

$$\text{type T: } \text{EUI} = \max (\text{XT} - \text{LT}, 900) \text{ seconds} \quad (151)$$

where the function **max** indicates the maximum of the values.

$$\text{type M: } \text{EUI} = \max (10\% \text{ of } (\text{LT} - \text{MT}), 900) \text{ seconds} \quad (152)$$

$$\text{type N: } \text{EUI} = 900 \text{ seconds} \quad (153)$$

In a preferred embodiment, the values 10% and 900 are parameters which can be set by an operator for the cache 110. In alternative embodiments, the 10% and 900 can be determined by the cache 110 responsive to global parameters of the cache 110. For example, the value 900 may be replaced by an average of EUI for all web objects 114 maintained in the cache 110.

The values shown herein are just initial estimates for EUI for each web document 114. In alternative embodiments, there might be many other ways to select initial estimates for EUI for each web document 114.

Similarly, in a preferred embodiment, the cache 110 estimates a value for EAI. In a preferred embodiment, the cache 110 makes an initial estimate for EAI equal to the mean value for EAI for all web documents 114 maintained in the cache 110.

At a step 212, the cache 110 determines which of the web objects 114 to refresh.

The cache 110 performs the step 212 by determining the probabilities  $\text{Psi}(t)$  and  $\text{Pri}(h)$  for each web object 114  $i$ , and selecting for refresh the web object 114  $i$  with the largest product  $\text{Pi}(\text{current time, current request-time})$ .

The cache 110 refreshes the selected web object 114 and repeats the step 212 so long as the cumulative sum of the products  $\text{Pi}(\text{current time, current request-time})$  is larger than a

1 selected threshold. As noted herein, in a preferred embodiment the selected threshold is between  
2 about 1% and about 5%, although a wide range of other values are likely also to be workable.

3

4 At a flow point 220, a selected refresh time (RT) for a particular web object 114  
5 is reached. The selected refresh time equals the latest load time (LT) plus EUI.

6

7 At a step 221, the cache 110 updates its estimated EUI for the web object 114, in  
8 response to an update history for the web object 114, as follows:

9

$$10 \quad \text{new EUI} = (1 - \alpha) (\text{old EUI}) + (\alpha) (UT - LT) \quad (171)$$

11

12 where UT = time the web object 114 is actually updated at the cache 110, and

13 LT = time the web object 114 is actually loaded or reloaded at the cache 110.

14

15 In a preferred embodiment, the value of the smoothing constant alpha is about  
16 0.4, but a wide range of values between 0 and 1 are likely to be usable.

17

18 In a preferred embodiment, EUI is only updated when the web object 114 is actu-  
19 ally updated at the cache 110. However, in alternative embodiments, EUI may be updated at  
20 other events, such as when the selected refresh time RT exceeds the sum (LT + EUI), that is, the  
21 web object 114 is not actually updated until after the estimated mean refresh time RT. More  
22 generally, in alternative embodiments, EUI may be updated responsive to (1) the last time at  
23 which the web object 114 was actually updated, (2) the amount of time actually passed since that  
24 update, and (3) any earlier estimate for EUI.

25

26 At a flow point 230, the particular web object 114 is requested by at least one cli-  
27 ent device 120.

28

29 At a step 231, the cache 110 updates its estimated EAI for the web object 114, in  
30 response to an request history for the web object 114, as follows:

31

$$32 \quad \text{new EAH} = (1 - \alpha) (\text{old EAH}) + (\alpha) (AH - LH) \quad (181)$$

33

34 where AH = request-time the web object 114 is actually requested by one of the client  
35 devices 120, and

1           LH = request-time the web object 114 was last requested by one of the client devices  
2           120.

3  
4           In a preferred embodiment, EAI is only updated when the web object 114 is actu-  
5 ally requested by one of the client devices 120. However, in alternative embodiments, EAI may  
6 be updated at other events, such as when the time between requests exceeds EAI, that is, the web  
7 object 114 is not actually requested until after the estimated mean request interval. More gener-  
8 ally, in alternative embodiments, EAI may be updated responsive to (1) the last time at which the  
9 web object 114 was actually requested, (2) the amount of request-time actually passed since that  
10 request, and (3) any earlier estimate for EAI.

11  
12           In a preferred embodiment, the value of the smoothing constant alpha is about  
13 0.4, but a wide range of values between 0 and 1 are likely usable.

14  
15           At a flow point 240, the cache 110 is ready to determine an estimated cache hit  
16 rate.

17  
18           At a step 241, the cache 110 sums the values of Pri (current request-time) for all  
19 web objects 114 in the cache 110.

20  
21           The sum of all Pri (current request-time) is an estimate of the probability that the  
22 next web object 114 to be requested is one of the web objects 114 maintained in the cache 110,  
23 that is, an estimate of the probability of a cache hit.

24  
25           At a flow point 250, the cache 110 is ready to select a particular web object 114  
26 for removal.

27  
28           At a step 251, the cache 110 determines which of the web objects 114 to remove.

29  
30           The cache 110 performs the step 213 by determining the load duration for the  
31 web object 114.

32  
33           
$$\text{new LD} = (1 - \alpha) (\text{old LD}) + (\alpha) (c + t) \quad (191)$$

34  
35           where alpha = a smoothing constant, preferably about 0.25.

1           c = actual connection time to the server device 130 for this load or reload. and

2           t = actual transmission time to the server device 130 for this load or reload.

3  
4           In a preferred embodiment, the value 0.25 for alpha is a parameter which can be  
5 set by an operator for the cache 110.

6  
7           In alternative embodiments, the value LD can be estimated by independently es-  
8 timating LDc and LDt, as follows:

9  
10           
$$LD = LDc + LDt \quad (192)$$

11  
12           where LDc = estimated connection time to the server device 130, and

13           LDt = estimated transfer time of the web object 114 from the server device 130.

14  
15           Each time the web object 114 is loaded or reloaded from the server device 130,  
16 the cache 110 revises its estimates for LDc and LDt, as follows:

17  
18           
$$\text{new LDc} = (1 - \alpha) (\text{old LDc}) + (\alpha) (c) \quad (193)$$

19  
20           
$$\text{new LDt} = (1 - \alpha) (\text{old LDt}) + (\alpha) (t) \quad (194)$$

21  
22           The cache 110 uses the estimated load duration LD for the web object 114, the  
23 size s of the web object 114, and the probability Pri (h) for each web object 114 i. and selects the  
24 web object 114 with the smallest product (LD/s) • Pri (current request) for removal.

25  
26           At a step 252, the cache 110 adds the value Pri (current request) it determined for  
27 the web object 114 to one of a set of summation buckets.

28  
29           In a preferred embodiment, each summation bucket is used to sum about 10,000  
30 sequential values of Pri (h). In an environment in which the mean size for web documents 114 is  
31 about eight kilobytes and each disk drive is about two gigabytes, each bucket therefore repre-  
32 sents the marginal value of about 1/25 of a disk drive.

33  
34           In a preferred embodiment, there are about 50 summation buckets. which are se-  
35 lected and filed in a round-robin manner. In an environment in which the mean size for web

1 documents 114 is about eight kilobytes and each disk drive is about two gigabytes. the et of 50  
2 buckets therefore represents the marginal value of about two disk drives.

3

4 At a step 253, the cache 110 deletes the web object 114 and its estimated values  
5 EUI and EAH.

6

7 *Alternative Embodiments*

8

9 Although preferred embodiments are disclosed herein, many variations are possi-  
10 ble which remain within the concept, scope, and spirit of the invention, and these variations  
11 would become clear to those skilled in the art after perusal of this application.

12



Claims

1  
2  
3  
4 1. A method of operating a cache, including the step of automatically re-  
5 freshing a set of objects maintained in said cache in response to an estimate for each particular  
6 object in said set that said object is stale.

7  
8 2. A method as in claim 1, wherein said step of automatically refreshing is  
9 responsive to an estimate for said particular object that said object will be requested soon.

10  
11 3. A method as in claim 1, wherein said step of automatically refreshing in-  
12 cludes the steps of  
13 for each said object, estimating a probability that said object will be next re-  
14 quested and will also be stale; and  
15 automatically refreshing a first object in response to said probability that said  
16 object will be next requested and will also be stale.

17  
18 4. A method as in claim 3, wherein said first object has a largest said prob-  
19 ability that said object will be next requested and will also be stale.

20  
21 5. A method as in claim 3, wherein said step of determining includes the  
22 steps of  
23 estimating a first probability distribution of client requests for said object;  
24 estimating a second probability distribution of server changes to said object; and  
25 estimating said probability that said object will be next requested and will also be  
26 stale in response to said first probability distribution and said second probability distribution.

27  
28 6. A method as in claim 3, wherein said step of estimating said probability  
29 that said object will be next requested and will also be stale is responsive to a product of said  
30 first probability distribution and said second probability distribution.

31  
32 7. A method as in claim 3, wherein said step of estimating includes the step  
33 of updating said estimated probability distribution for each object in response to a request his-  
34 tory for said object.

1                   8.     A method as in claim 7, wherein said step of updating includes the steps  
2     of  
3                   determining an initial estimated probability distribution of client requests for said  
4     object; and  
5                   updating said estimated probability distribution of client requests for said object  
6     in response to said request history.

7  
8                   9.     A method as in claim 7, wherein said step of updating includes the steps  
9     of  
10                  determining an initial estimated probability distribution of client requests for said  
11     object;  
12                  determining a new estimated probability distribution of client requests for said  
13     object in response to said request history; and  
14                  determining a composite value of said initial estimated probability distribution  
15     and said new estimated probability distribution.

16  
17                  10.    A method as in claim 3, wherein said step of determining includes the step  
18     of updating said probability for each object in response to an update history for said object.

19  
20                  11.    A method as in claim 10, wherein said step of updating includes the steps  
21     of  
22                  determining an initial estimated probability distribution of server changes to said  
23     object; and  
24                  updating said estimated probability distribution of server changes to said object in  
25     response to said update history.

26  
27                  12.    A method as in claim 10, wherein said step of updating includes the step  
28     of determining an initial estimated probability distribution of server changes to said object in  
29     response to information provided with said object by a server device.

30  
31                  13.    A method as in claim 10, wherein said step of updating includes the steps  
32     of  
33                  determining an initial estimated probability distribution of server changes to said  
34     object;

1                   determining a new estimated probability distribution of server changes to said  
2 object in response to said update history; and  
3                   determining a composite value of said initial estimated probability distribution  
4 and said new estimated probability distribution.

5  
6                   14. A method of operating a cache, including the step of automatically re-  
7 freshing a set of objects maintained in said cache in response to an estimate for each particular  
8 object in said set that said object will be requested soon.

9  
10                  15. A method as in claim 14, wherein said step of automatically refreshing in-  
11 cludes the steps of  
12                   for each said object, estimating a probability that said object will be next re-  
13 quested and will also be stale; and  
14                   automatically refreshing a first object in response to said probability that said  
15 object will be next requested and will also be stale.

16  
17                  16. A method of operating a cache, including the steps of  
18 Determining, for each object in a set of objects maintained in said cache, a probability of client  
19 requests for said object and an expected load duration for said object; and  
20                   Selecting a particular object for removal in response to said probability of client  
21 requests for said object and said expected load duration for said object.

22  
23                  17. A method as in claim 16, wherein said step of selecting is responsive to a  
24 product of load duration per unit size and said probability of client requests for each said object.

25  
26                  18. A method as in claim 16, wherein said step of determining includes the  
27 step of updating said probability for each object in response to a request history for said object.

28  
29                  19. A method as in claim 18, wherein said step of updating includes the steps  
30 of  
31                   determining an initial estimated probability distribution of client requests for said  
32 object; and  
33                   updating said estimated probability distribution of client requests for said object  
34 in response to said request history.

35

1                   20.    A method as in claim 18, wherein said step of updating includes of steps  
2    of  
3                   determining an initial estimated probability distribution of client requests for said  
4    object;  
5                   determining a new estimated probability distribution of client requests for said  
6    object in response to said request history; and  
7                   determining a composite value of said initial estimated probability distribution  
8    and said new estimated probability distribution.

9  
10                  21.    A method as in claim 16, wherein said step of determining includes the  
11   step of updating said expected load duration for each object in response to a history for said ob-  
12   ject.

13  
14                  22.    A method as in claim 21, wherein said step of updating includes the steps  
15   of  
16                  determining an initial expected load duration for said object; and  
17                  updating said expected load duration for said object in response to said request  
18   history.

19                  23.    A method as in claim 21, wherein said step of updating includes the steps  
20   of  
21                  determining an initial expected load duration for said object;  
22                  determining a new expected load duration for said object in response to said his-  
23   tory; and  
24                  determining a composite value of said initial expected load duration and said new  
25   expected load duration.

26  
27                  24.    A method as in claim 21, wherein said step of updating includes the step  
28   of updating said expected load duration for a particular object when said object reloaded.

29  
30                  25.    A method of operating a cache, including the steps of  
31                  for a plurality of objects maintained in said cache, determining an estimated  
32   probability distribution of cliente requests for said object; and  
33                  cumulating an element of said estimated probability distribution for a set of said  
34   objects.

35

1                   26.    A method as in claim 25, wherein said set of said objects includes a set of  
2 recently deleted objects.

3  
4                   27.    A system, including  
5 a cache; and  
6 means for automatically refreshing a set of objects maintained in said cache in re-  
7 sponse to an estimate for each particular object in said set that said object is stale.

8  
9                   28.    A system as in claim 27, wherein said means for automatically refreshing  
10 is responsive to an estimate for said particular object that said object will be requested soon.

11  
12                   29.    A system as in claim 27, wherein said means for automatically refreshing  
13 includes  
14 means for estimating a probability, for each said object, that said object will be  
15 next requested and will also be stale; and  
16 means for automatically refreshing a first object in response to said probability  
17 that said object will be next requested and will also be stale.

18  
19                   30.    A system as in claim 29, wherein said first object has a largest said prob-  
20 ability that said object will be next requested and will also be stale.

21  
22                   31.    A system as in claim 29, wherein said means for determining includes  
23 memory storing an estimate of a first probability distribution of client requests for  
24 said object;  
25 memory storing an estimate of a second probability distribution of server changes  
26 to said object; and  
27 means for estimating said probability that said object will be next requested and  
28 will also be stale in response to said first probability distribution and said second probability  
29 distribution.

30  
31                   32.    A system as in claim 29, wherein said means for estimating said probabil-  
32 ity that said object will be next requested and will also be stale is responsive to a product of said  
33 first probability distribution and said second probability distribution.

1                   33. A system as in claim 29, wherein said means for estimating includes  
2 means for updating said estimated probability distribution for each object in response to a re-  
3 quest history for said object.

4  
5                   34. A system as in claim 33, wherein said means for updating includes  
6 memory storing an initial estimated probability distributio of client re-  
7 quests for said; and  
8 means for updating said estimated probability distribution of client re-  
9 quests for said object in response to said request history.

10  
11                  35. A system as in claim 33, wherein said means for updating includes means  
12 for estimating an initial estimated probability distribution of client requests for said object;  
13 means for estimating a new estimated probability distribution of client requests  
14 for said object in response to said request history; and  
15 means for determining a composite value of said initial estimated probability dis-  
16 tribution and said new estimated probability distribution.

17  
18                  36. A system as in claim 29, wherein said means for determining includes  
19 means for updating said probability for each object in response to an update history for said ob-  
20 ject.

21                  37. A system as in claim 36, wherein said means for updating includes  
22 Means for estimating an initial estimated probability distribution of server  
23 changes to said object; and  
24 means for updating said estimated probability distribution of server changes to  
25 said object in response to said update history.

26  
27                  38. A system as in claim 36, wherein said means for updating includes means  
28 for estimating an initial estimated probability distribution of server changes to said object in re-  
29 sponse to information provided with said object by a server device.

30  
31                  39. A system as in claim 36, wherein said means for updating includes  
32 means for estimating an initial estimated probability distribution of server  
33 changes to said object;

34 means for estimating a new estimated probability distribution of server changes to  
35 said object in response to said update history; and

1 means for determining a composite value of said initial estimated probability dis-  
2 tribution and said new estimated probability distribution.

3  
4 40. A system, including

5 a cache; and

6 means for automatically refreshing a set of objects maintained in said cache in re-  
7 sponse to an estimate for each particular object in said set that said object will be requested soon.

8  
9 41. A system as in claim 40, wherein said means for automatically refreshing  
10 includes the steps of

11 for each said object, estimating a probability that said object will be next re-  
12 quested and will also be stale; and

13 automatically refreshing a first object in response to said probability that said  
14 object will be next requested and will also be stale.

15  
16 42. A system for operating a cache, including

17 means for determining, for each object in a set of objects maintained in said  
18 cache, a probability of client requests for said object and an expected load duration for said ob-  
19 ject; and

20 means for selecting a particular object for removal in response to said probability  
21 of client requests for said object and said expected load duration for said object.

22  
23 43. A system as in claim 42, wherein said means for selecting is responsive to  
24 a product of load duration per unit size and said probability of client requests for each said ob-  
25 ject.

26  
27 44. A system as in claim 42, wherein said means for determining includes  
28 means for updating said probability for each object in response to a request history for said ob-  
29 ject.

30  
31 45. A system as in claim 44, wherein said means for updating includes means  
32 for estimating an initial estimated probability distribution of client requests for said object; and  
33 means for updating said estimated probability distribution of client requests for  
34 said object in response to said request history.

1                   46.    A system as in claim 44, wherein said means for updating includes means  
2   for estimating an initial estimated probability distribution of client requests for said object;  
3                   means for estimating a new estimated probability distribution of client requests  
4   for said object in response to said request history; and  
5                   means for determining a composite value of said initial estimated probability dis-  
6   tribution and said new estimated probability distribution.

7  
8                   47.    A system as in claim 42, wherein said means for determining includes  
9   means for updating said expected load duration for each object in response to a history for said  
10   object.

11  
12                  48.    A system as in claim 47, wherein said means for updating includes means  
13   for estimating an initial expected load duration for said object; and  
14                  means for updating said expected load duration for said object in response to said  
15   request history.

16  
17                  49.    A system as in claim 47, wherein said means for updating includes means  
18   for estimating an initial expected load duration for said object;  
19                  means for estimating a new expected load duration for said object in response to  
20   said history; and  
21                  means for determining a composite value of said initial expected load duration  
22   and said new expected load duration.

23  
24                  50.    A system as in claim 47, wherein said means for updating includes means  
25   for updating said expected load duration for a particular object when said object is reloaded.

26  
27                  51.    A system for operating a cache, including means for determining, for each  
28   object in a set of objects maintained in said cache, an estimated probability distribution of client  
29   requests for said object; and  
30                  means for cumulating an element of said estimated probability distribution for a  
31   set of said objects.

32  
33                  52.    A system as in claim 51, wherein said set of said objects includes a set of  
34   recently deleted objects.



1/2




FIG. 1

2/2

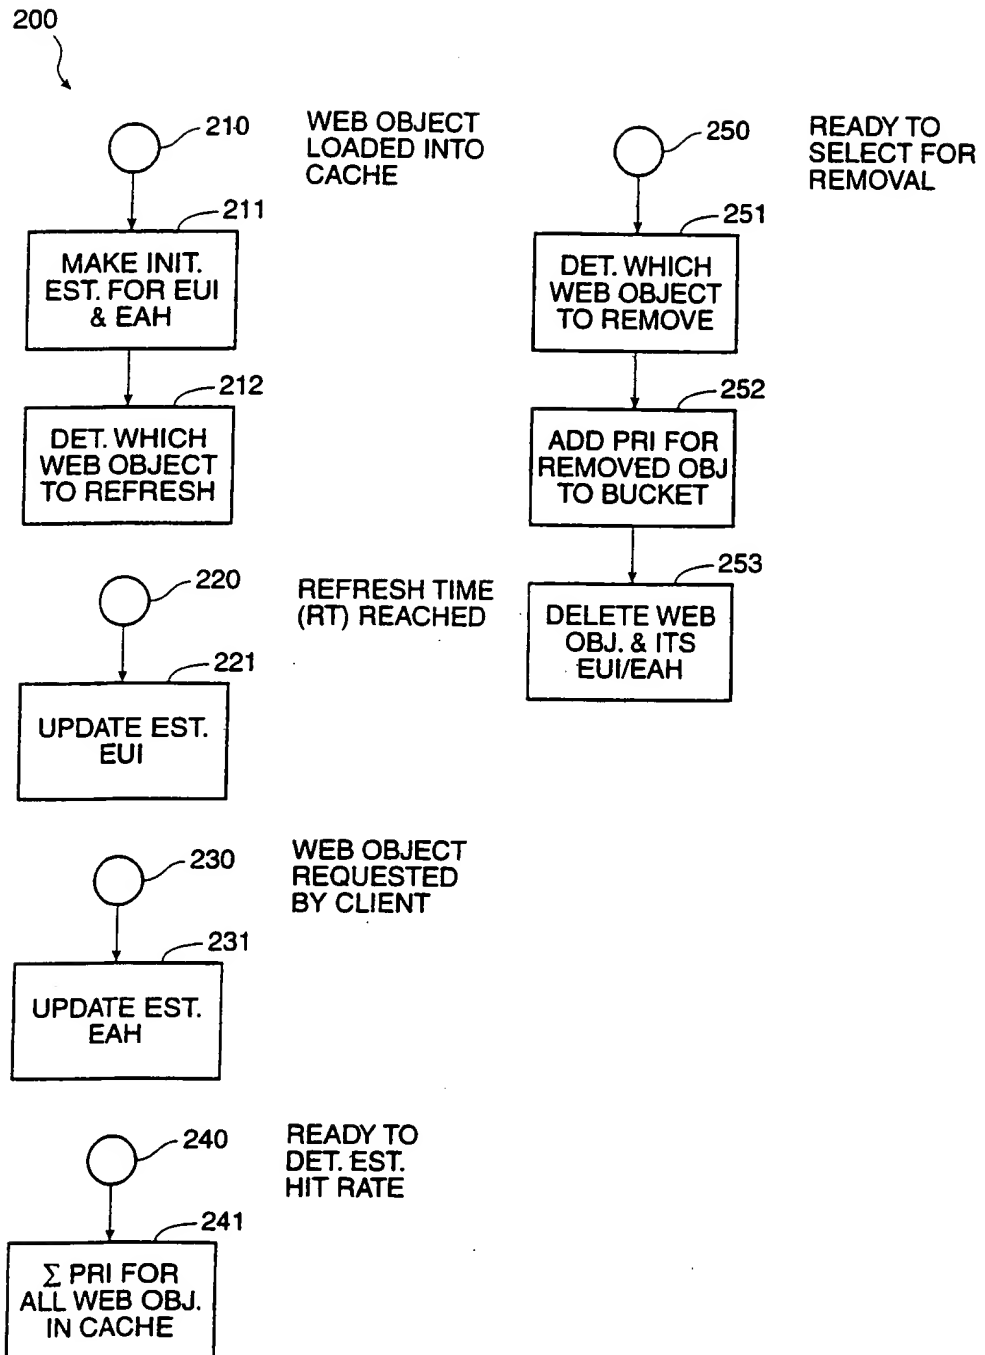


FIG. 2

# INTERNATIONAL SEARCH REPORT

Inte	Application No
PCT/US	98/20719

**A. CLASSIFICATION OF SUBJECT MATTER**  
 IPC 6 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
 IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X          A	DIAS G. ET AL: "A Smart Internet Caching System" PROCEEDINGS OF THE INET'96 CONFERENCE, MONTREAL, CANADA, 24 - 28 June 1996, XP002086721 htt://www.isoc.org/inet96/proceedings/a4/a4_3.htm see the whole document  <div style="text-align: center;">---</div> <div style="text-align: center;">-/--</div>	1-3, 5, 14, 15, 27-29, 31, 40, 41           4, 6-13, 16-26, 30, 32-39, 42-52

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

\* Special categories of cited documents :

- \* "A" document defining the general state of the art which is not considered to be of particular relevance
- \* "E" earlier document but published on or after the international filing date
- \* "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- \* "O" document referring to an oral disclosure, use, exhibition or other means
- \* "P" document published prior to the international filing date but later than the priority date claimed

- \* "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- \* "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- \* "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- \* "&" document member of the same patent family

Date of the actual completion of the international search

4 December 1998

Date of mailing of the international search report

21/12/1998

Name and mailing address of the ISA  
 European Patent Office, P.B. 5818 Patentlaan 2  
 NL - 2280 HV Rijswijk  
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.  
 Fax: (+31-70) 340-3016

Authorized officer

Fournier, C

# INTERNATIONAL SEARCH REPORT

Inter Application No  
PCT/US 98/20719

## C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	GLASSMAN S: "A caching relay for the World Wide Web" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 27, no. 2, November 1994, page 165-173 XP004037987 see page 166, right-hand column, line 23 - page 167, right-hand column, line 3 see page 172, right-hand column, paragraph 11 - page 173, left-hand column, paragraph 12	1,2,14, 27,28,40
A	----- ZHIMEI JIANG ET AL: "Prefetching links on the WWW" 1997 IEEE INTERNATIONAL CONFERENCE ON COMMUNICATIONS. TOWARDS THE KNOWLEDGE MILLENNIUM. ICC '97. CONFERENCE RECORD (CAT. NO.97CH36067), PROCEEDINGS OF ICC'97 - INTERNATIONAL CONFERENCE ON COMMUNICATIONS, MONTREAL, QUE., CANADA, 8-12 JUNE 1997, pages 483-489 vol.1, XP002086568 ISBN 0-7803-3925-8, 1997, New York, NY, USA, IEEE, USA see page 483, left-hand column, paragraph 1 - page 484, right-hand column, paragraph 3A see page 4898, left-hand column, paragraph 6	1-52
A	----- DINGLE A ET AL: "Web cache coherence" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 28, no. 11, May 1996, page 907-920 XP004018195 see page 907, left-hand column, paragraph 1 - page 911, right-hand column, paragraph 2.5 see page 917, left-hand column, paragraph 4 - page 918, left-hand column, paragraph 4.2	1-52
A	----- NABESHIMA M: "The Japan Cache Project: an experiment on domain cache" COMPUTER NETWORKS AND ISDN SYSTEMS, vol. 29, no. 8-13, September 1997, page 987-995 XP004095297 see page 991, left-hand column, paragraph 4.2 - page 992, right-hand column, paragraph 4.2.5; figure 2 -----	1-52